Social Choice, Fairness, and Learning

- Anne-Marie George -

BSc at University Bonn & MSc at TU Berlin in **Mathematics**

- Focus on Discrete Mathematics and Optimisation

PhD at University College Cork in **Computer Science**

- Preference Inference Based on Lexicographic and Pareto Models

Single User's Preferences Inference & Consistency

tdoc at TU Berlin in Efficient Algorithms Group

tational Social Choice and in particular **Voting Theory**

Multi. Agent's Preferences Aggregation

Postdoc a ASR

- **Social Ch airness** and **Machine/Reinforcement Learning**

Multi. Agent's Preferences Learning & Elicitation

Computational Social Choice Theory

- At the intersection of Economics, Computer Science and Mathematics
- Analysing mechanisms to make (fair) group decisions, e.g.,
  - matchings (kidney donation schemes, housing markets,…)
  - fair division (splitting divisible or indivisible goods)
  - voting (analysing voting schemes: fairness, manipulability, …)
  - …
- All problems involve **agents with preferences** (e.g. rankings over alternatives)
- Rational agents – close to Game Theory

Defining Fairness of an Allocation Policy

- Policy: Probability of selecting (sets of) agents
  Utility: Based on selection.
- Meritocracy: One-time selection
  Merit of selecting an agent = expected marginal contribution to utility
  T. Kleine Buening, M. Segal, D. Basu, A. George, C. Dimitrakakis.
  On Meritocracy in Optimal Set Selection. EAAMO'22
- Policy Fairness: Dynamic (via MDP formulation):
  Selection changes Societal Bias/Fairness State
  Weighted sum of "fairness" in each reached state ≈ reward function
  M. Segal, A. George, C. Dimitrakakis.
  Policy Fairness and Unknown Bias Dynamics in Sequential Allocations. EAAMO'23

Modelling Feedback in Society with MDPs

- Feedback: Of e.g. affirmative actions, recourse, …
- May change preferences or distribution of population of agents
- Questions:
  Long term effect?
  Validity of recourse?
  Fairness over time?

Setting: College Admissions

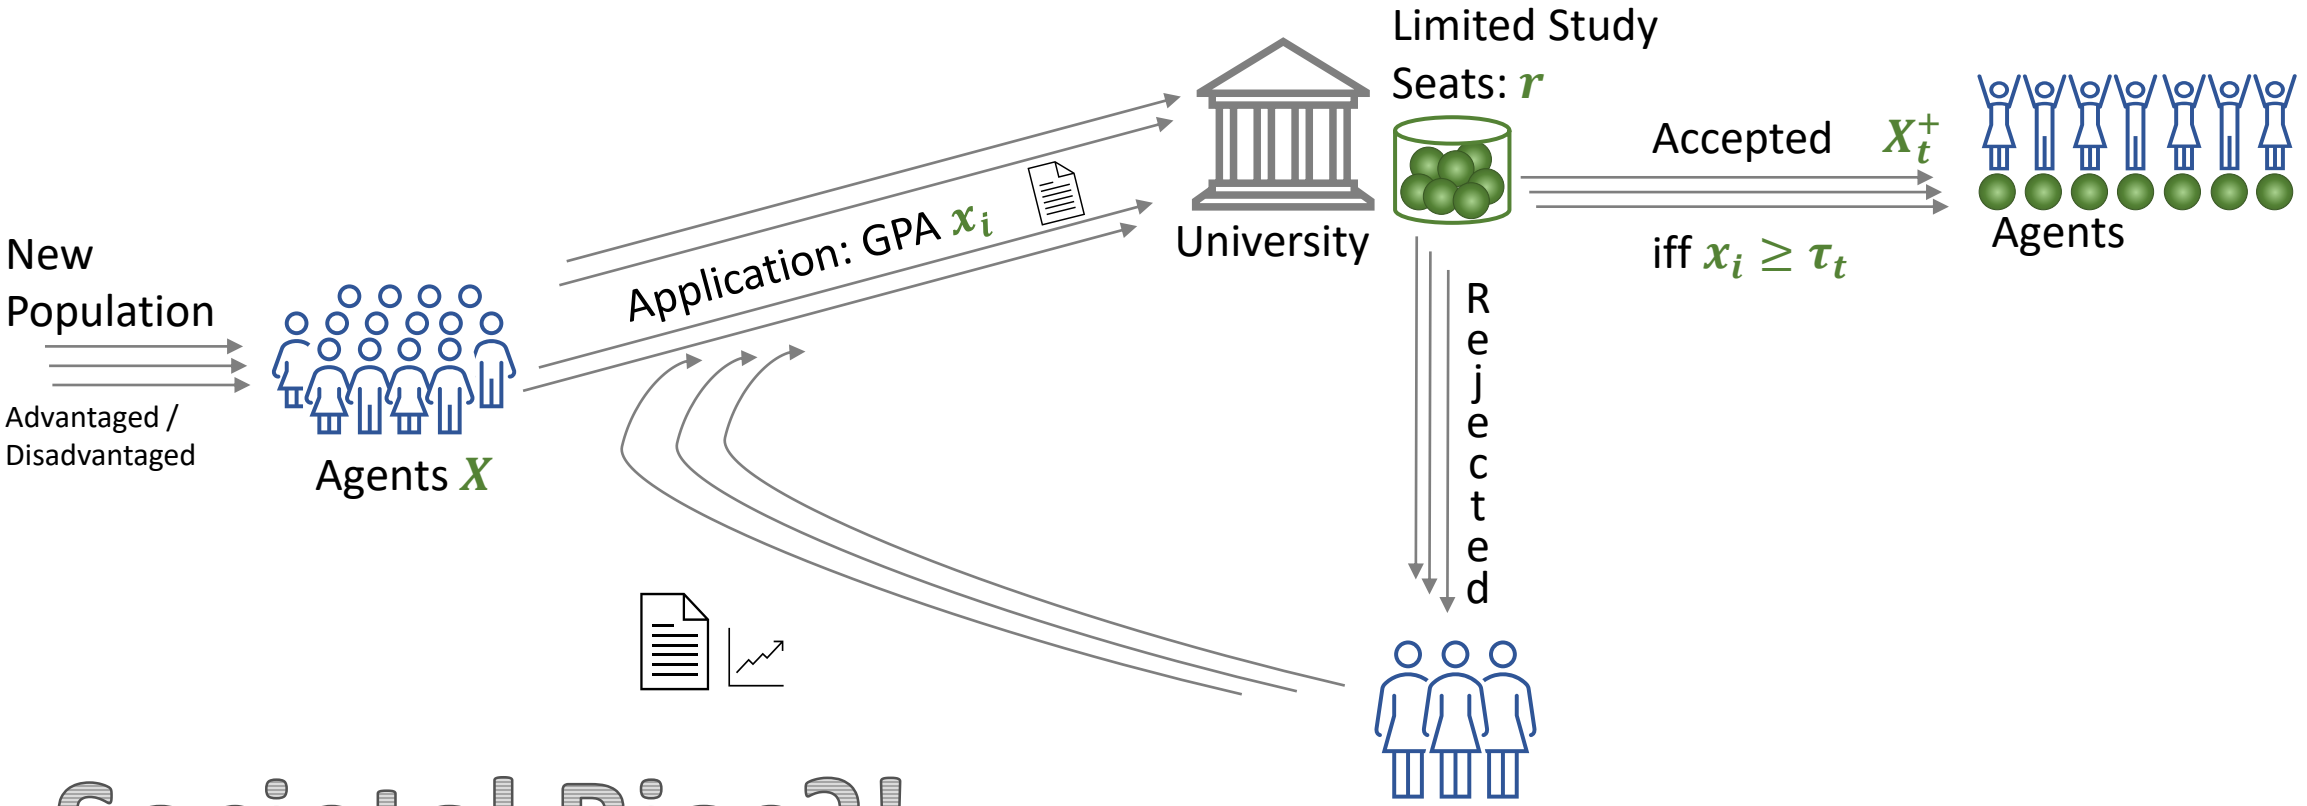

Societal Bias?!

Increase? (reinforcing stereotypes, …)
Decrease? (better representation, …)
Levels out over time?

Setting: College Admissions

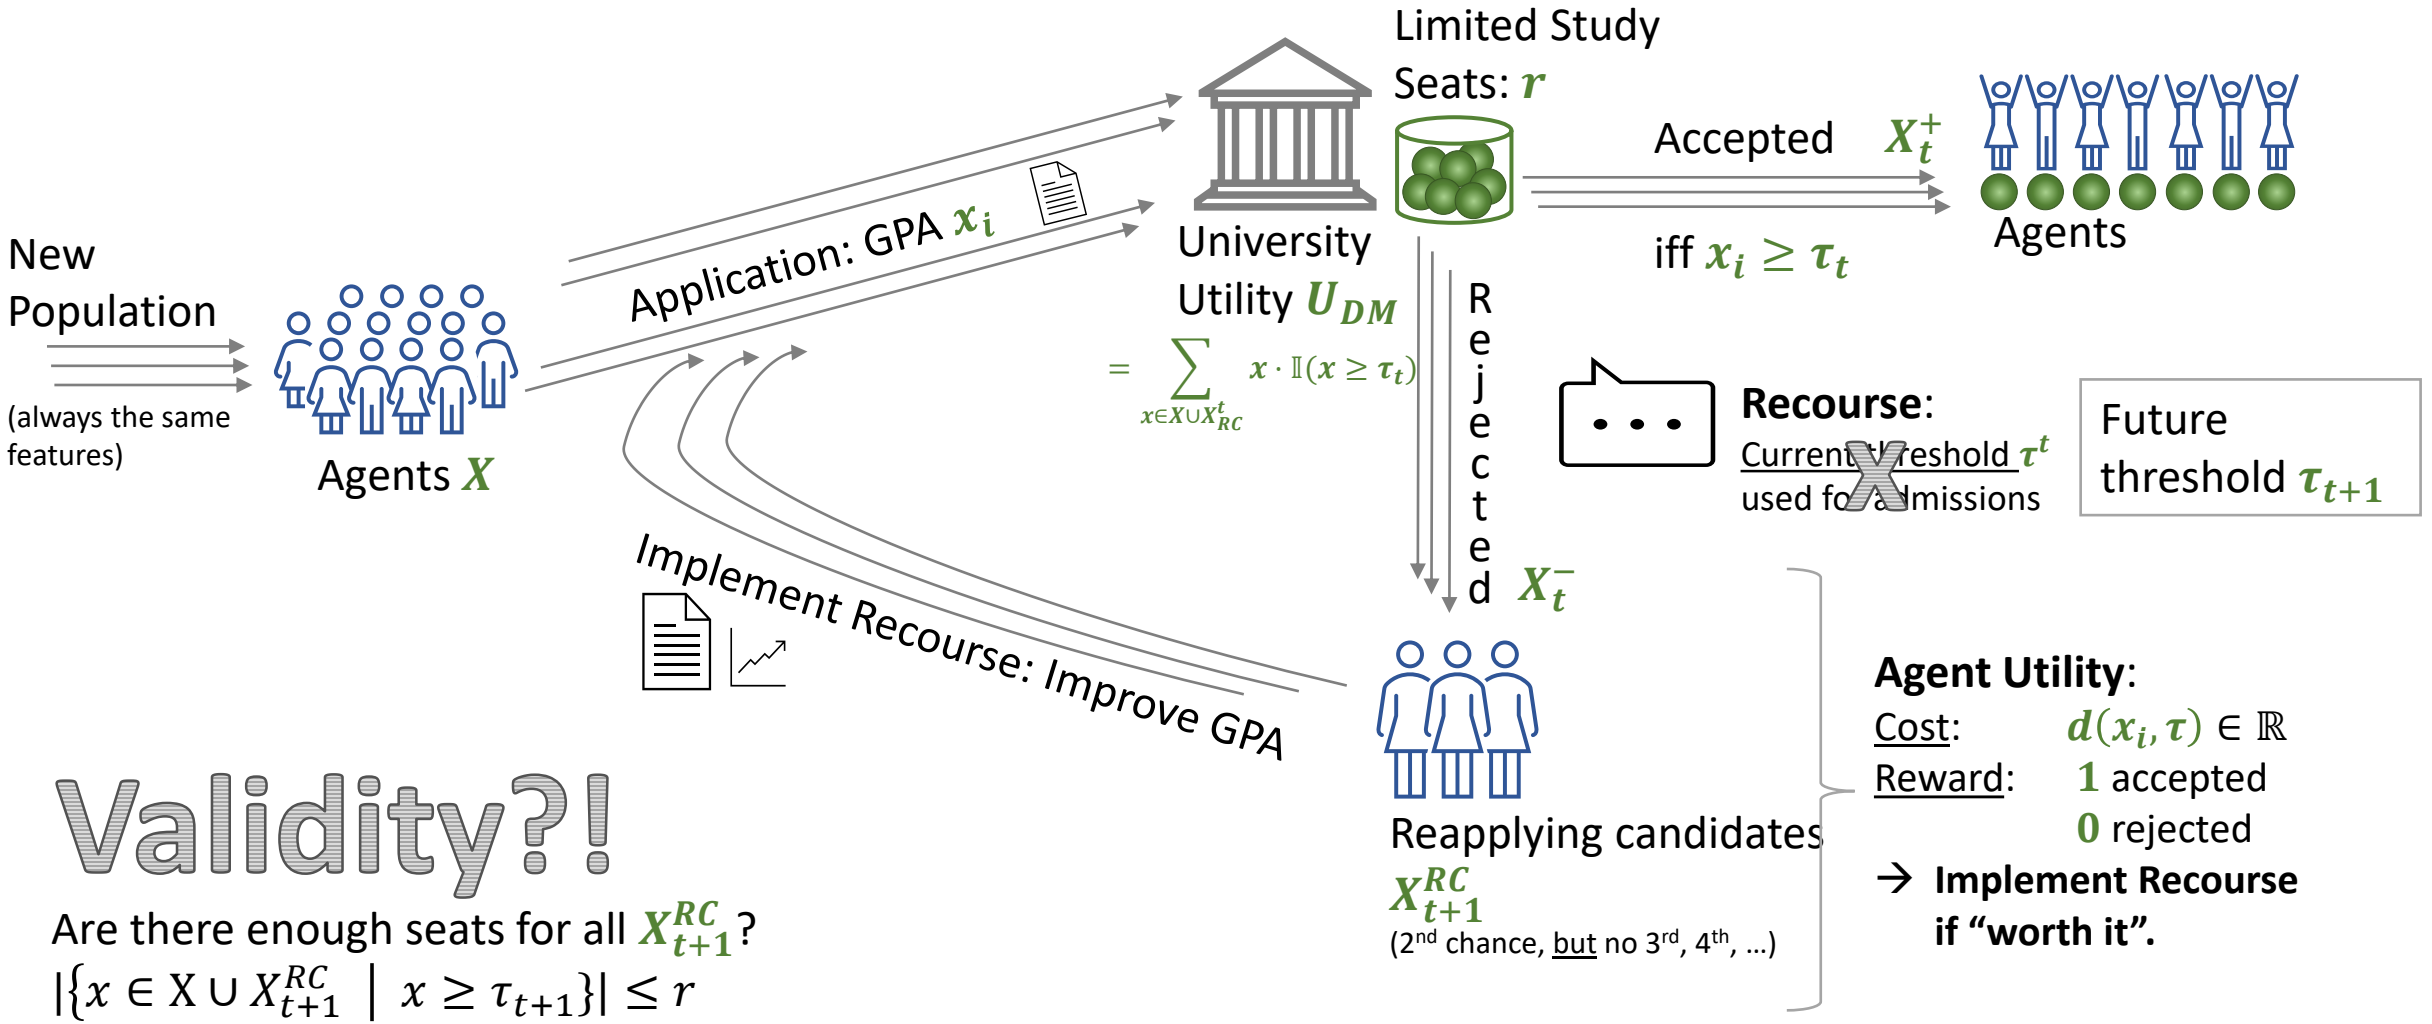

Validity?!

Are there enough seats for all X^{RC}_{t+1}?

$|\{x \in X \cup X^{RC}_{t+1} \mid x \geq \tau_{t+1}\}| \leq r$

Learning Stable Matchings

A. Athanasopoulos, A. George, C. Dimitrakakis. Probably Correct Optimal Stable Matching for Two-Sided Market Under Uncertainty. AAMAS'25.

At every time step: Match agents to jobs (or sth. to sth.else)

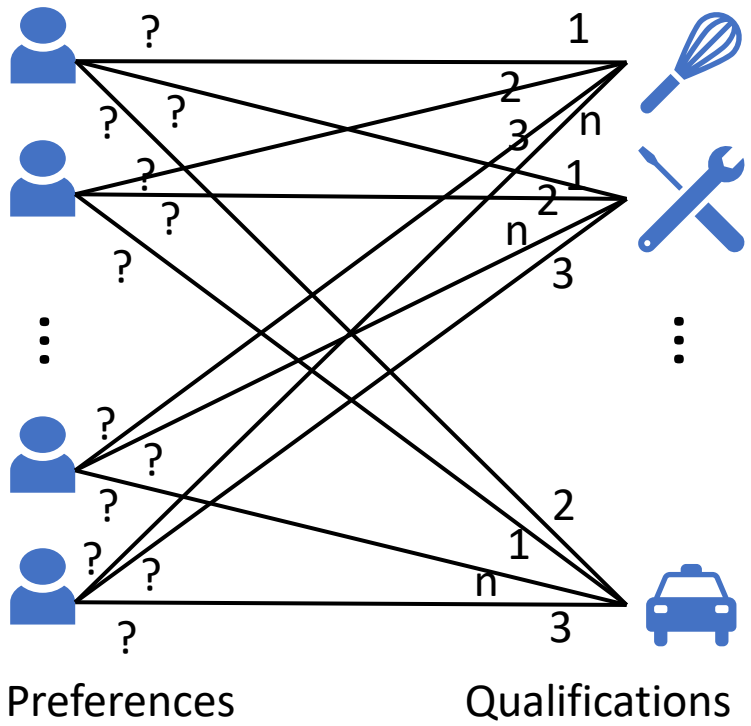

Learning Stable Matchings

A. Athanasopoulos, A. George, C. Dimitrakakis. Probably Correct Optimal Stable Matching for Two-Sided Market Under Uncertainty. AAMAS'25.

At every time step: Match agents to jobs (or sth. to sth.else)

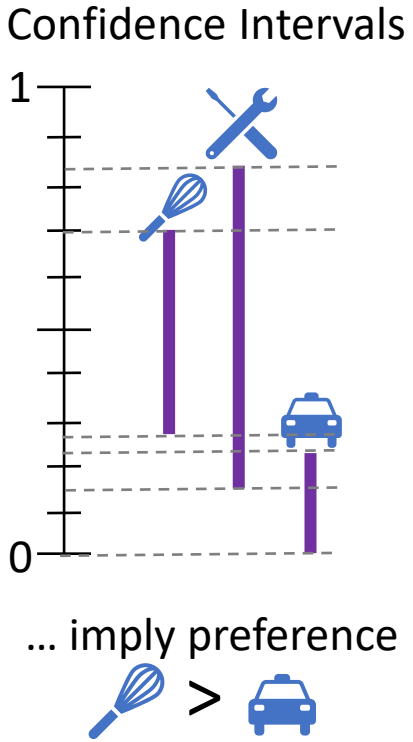

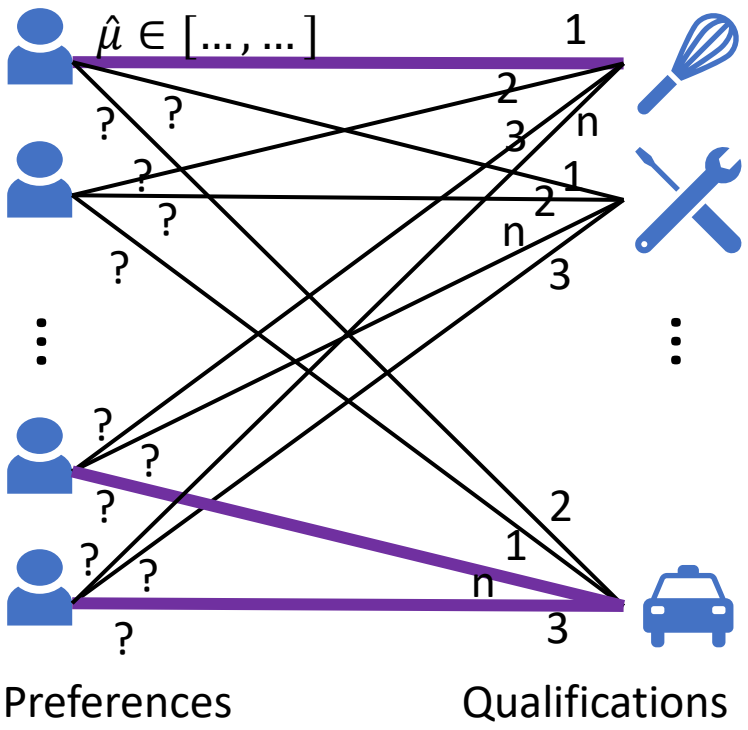

Use Bandit style Sampling to Learn an Optimal Matching (Pure Exploration Setting)

- Uniform Sampling
- UCB
- Adaptive Based on Properties of Matchings

Eliciting Votes with Dueling Bandits

A. George, C. Dimitrakakis. Eliciting Kemeny Rankings. AAAI'24.

$C = \{1, \dots, k\}$ candidates

- V voters with preference rankings

- <u>Goal:</u> Aggregate the preferences - find a Kemeny ranking!
- <u>Problem</u>: What if voters can't specify their whole preference ranking?
  - Cognitive effort, Number of candidates too large
- <u>Idea</u>: Only **ask for few comparisons**, e.g. "$a \succ b$?", and try to **approximate the voting outcome.**

Dueling Bandits Framework – Sampling Q

$C = \{1, \dots, k\}$ candidates (arms)

- Choose a **random (uar)** voter in V and ask some pairwise comparison (i, j)

- Outcome of pulling arms i and j:

winning probabilities

- $i \succ j$: with prob. $q_{ij} = \frac{|\{v \in V \mid i \succ_v j\}|}{|V|}$
- $j \succ i$: with prob. $q_{ji} = 1 - q_{ij}$

The "*dueling bandits*"

arm i vs. arm j

PAC Sample Complexity

- Approximation guarantee based on confidence intervals
  → Make use of properties of preferences (transitivity, …)
  → Dependent on voting rule

- Concentration inequalities for guarantee of probability

- Compute necessary number of samples for PAC bounds

Other Research

- Integreat:
  - PI for Explainability & Fairness
  - PhD project: Handling Inconsistencies (particularly for Preference Learning)
- RCN Researcher Project for Early Career Scientists (FRIPRO)
  - **So**cietal **P**reference **H**andling and **E**licitation
    - pending resubmission with adjusted budget: 1 PhD, 1 Postdoc
- Other
  - Starterpakke:1 PhD

Teaching

- Before:
  - Reinforcement learning (IN-STK5100/9100) (Vår 2021-23)
  - Adaptive Databased Decision Making (IN-STK5000/9000) (Høst 2021-22)
- Now:
  - Logiske Metoder (Vår 2025)
- Future: ???
  - Probablistic ML, Bachelors Course (Høst 2026 ???)
  - RL, Masters Course ?
  - Multi-Agent Systems, Masters Course?